

Crime, Punishment, and Evolution in an Adversarial Game*

Michael McBride[†]

Department of Economics
University of California, Irvine

Ryan Kendall

Department of Economics
University of California, Irvine

Martin B. Short

Department of Mathematics
University of California, Los Angeles

Maria R. D’Orsogna

Department of Mathematics
California State University, Northridge

This version: September 13, 2012

Abstract

We examine the game theoretic properties of a model of crime first introduced by Short, Brantingham, and D’Orsogna (Short et al. 2010) as the SBD Adversarial Game. We identify the rationalizable strategies and one-shot equilibria under multiple equilibrium refinements. We further show that SBD’s main result about the effectiveness of defecting-punishers in driving the system to evolve to the cooperative equilibrium under an imitation dynamic does generalize to a best response dynamic, although the nature of this strategy’s role differs significantly between the two dynamics. The analysis reveals that the positive externality in punishing crime in the SBD game converts the adversarial setting from a social dilemma to a coordination game. We provide policy implications and lessons learned about the evolution of cooperation more generally.

JEL Classifications: C73, D74.

Keywords: cooperation, punishment, evolution.

*The material herein is based upon work supported by the Army Research Office Award No. #W911NF-11-1-0332. McBride and Short were also supported by the Air Force Office of Scientific Research Award No. FA9550-10-1-0569.

[†]Corresponding author. 3151 Social Science Plaza, Irvine, CA, 92697-5100, mcbride@uci.edu.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 13 SEP 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Crime, Punishment, and Evolution in an Adversarial Game		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California Los Angeles, Department of Mathematics, Los Angeles, CA, 90095		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We examine the game theoretic properties of a model of crime first introduced by Short, Brantingham, and D'Orsogna (Short et al. 2010) as the SBD Adversarial Game. We identify the rationalizable strategies and one-shot equilibria under multiple equilibrium refinements. We further show that SBD's main result about the effectiveness of defecting-punishers in driving the system to evolve to the cooperative equilibrium under an imitation dynamic does generalize to a best response dynamic although the nature of this strategy's role differs significantly between the two dynamics. The analysis reveals that the positive externality in punishing crime in the SBD game converts the adversarial setting from a social dilemma to a coordination game. We provide policy implications and lessons learned about the evolution of cooperation more generally.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 29	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1 Introduction

While various mechanisms can sustain cooperative behavior in social dilemmas, numerous theoretical models, case studies, and experimental data have shown the punishment of defectors to be particularly effective (e.g., Ostrom 1990; Ledyard 1995; Fehr and Gächter 2000; Henrich and Boyd 2001; Cinyabuguma, Page, and Putterman 2005; Takahashi 2010; Boyd, Gintis, and Bowles 2010; Fudenberg and Pathak 2010; Chaudhuri 2011; Xiao and Hauser 2011). Given that the punishment of defectors is itself a second order social dilemma, recent attention has been given to how cooperative and punishing behavior coevolve in evolutionary settings (e.g., Hauert et al. 2007; Carpenter 2007; Levine and Pesendorfer 2007; Fehr and Fischbacher 2004; Bowles and Gintis 2011). Most of these works examine the emergence of cooperative and punishing behaviors in the standard social dilemma games, such as the Prisoner’s Dilemma, Public Good Games, and Common Property Resource Games. The lack of cooperation in these settings is driven by the tension between personal and collective interests: players maximize their personal interest which, in turn, has an indirect and negative effect on the collective interest.

Standard social dilemmas pervade much of social and political life, but individuals in many societies—especially those with weak governments or state capacity—face dilemmas that are more adversarial in nature than standard social dilemmas. Such settings have been largely ignored in the literature, yet an important exception is Short, Brantingham, and D’Orsogna (2010), hereafter SBD. In their work, the authors introduce an evolutionary game in which two actors are selected at random, one placed in a potential criminal role, the other placed in a potential victim role. The former decides whether or not to steal from the latter, and the latter decides whether or not to report this theft to authorities, if it takes place. Whether a report leads to conviction, which directly affects the expected net benefit of reporting a crime, is probabilistically increasing in the population’s overall proclivity to cooperate with authorities (i.e., proclivity to punish criminals). SBD modify the traditional cooperate and defect strategies to account for punishing behavior, and then show that under

a specific imitation dynamic, the presence of a novel strategy type—dubbed the Informant—that both commits crimes and cooperates with authorities, is highly influential in driving the system away from the “Dystopian” state of high crime and toward the efficient, no-crime, “Utopian” steady state. Specifically, the presence of Informants is a sufficient, but not necessary, condition for achieving the Utopian state.

Our paper more closely examines the game theoretic properties of the SBD Adversarial Game. We consider the Adversarial Game worthy of study for multiple reasons. It provides an original depiction of an adversarial setting but also presents a simple formalization of the punishment of defectors (criminals) that explicitly captures the positive externality in punishment present in many societies: civilians play a crucial role in the self-regulation of pro-social norms (e.g., Sampson and Groves 1989; Skogan 1990; Bursik and Grasmick 1993), but fear of retaliation may lead to disengagement from law enforcement and the proliferation of criminal behavior (e.g., Gambetta 1988; Beittel 2009). Furthermore, the SBD model uses a simple, stylized schema of strategy types, including the novel Informant type, that is of independent interest. Finally, though the game’s design is directly influenced by features of crime and punishment in disorganized societies, the model’s simple structure has the potential to illuminate our understanding of the emergence of cooperative behavior in a more general way.

In this paper, we recast the SBD Adversarial Game in a classic game theoretic setting. A specific goal is to assess the robustness of SBD’s main finding regarding the influential role of Informants in the evolution of cooperation. From a broader perspective, we would like to identify what the Adversarial Game can teach us about the evolution of cooperation. As part of that goal we seek to identify exactly how the SBD game compares with other, more commonly studied social dilemma games. We thus present both static and evolutionary analyses. We first examine the one-shot Adversarial Game and fully characterize the set of rationalizable strategies and (Bayesian) Nash Equilibria. As a bridge between the static and evolutionary analyses, we then identify which equilibria survive two equilibrium refinements:

Evolutionary Stable Strategy and Trembling Hand Perfect Equilibrium. We finally turn to evolutionary analysis. Whereas SBD assume a particular imitation dynamic, we examine the evolutionary path of cooperative behavior under a simple best response dynamic.

Two main findings and two main lessons emerge from the analysis. Unlike other one-shot social dilemma games in which the unique equilibrium is inefficient, the one-shot Adversarial Game has multiple equilibria, one of which is efficient – Utopia. In effect, the specific form of the positive externality in reporting (i.e., punishing) converts the second order public goods problem into a second order coordination game, which in turn converts the overall game into a non-standard coordination game. A key lesson is that an institution that converts the second order public goods problem from a social dilemma into a coordination game can foster crime-detering punishment. We also show that SBD’s main result regarding the power of Informants to foster a low-crime society generalizes to the case of best response dynamics, though in a much different way than in the imitation dynamics of SBD. That is, we find the availability of the Informant strategy to be a necessary, but insufficient, condition to maintain a low-crime state, which will only occur given appropriate parameters and initial conditions. This finding suggests a modification to SBD’s policy recommendation of converting defectors into Informants to improve overall cooperation, under best response dynamics. Our analysis identifies settings under which this policy is more likely or less likely to work.

We emphasize that this Adversarial Game does not manifest the second-order free-rider problem in which the collective level of punishment on defectors is itself a type of public good problem, as in Fehr and Gächter (2000). Rather, the second-order punishment game is a coordination game in which the expected net benefits of punishing are increasing in the overall proportion of others that also punish. Thus, this Adversarial Game is closer in spirit to social dilemma models, such as Ostrom, Walker, and Gardner (1992) and Boyd, Gintis, and Bowles (2010), in which punishers must solve a coordination problem. Our findings thus complement prior theoretical work by showing how coordinated punishments can enforce cooperative behavior in a directly adversarial setting.

Our work also complements prior work in the economics literature on crime. In his review of literature, Ehrlich (1996) distinguishes two main strands of recent work: market models of crime and optimal crime control policies. Our paper resides in the former, though we do offer suggestions regarding the latter. Prior models have considered the equilibrium patterns of crime in static settings (e.g., Fender 1999), while others have considered the instability of equilibria (e.g., Eaton and Wen 2008). Our paper differs by considering an stylized, evolutionary setting, and looking at the evolution and steady states of crime and reporting. In this way we draw more direct comparisons to the evolution of cooperation literature.

2 The Adversarial Game

Consider a population of expected-payoff maximizing actors, $i = 1, \dots, N$, where N is large and even. Each actor has an endowment of 1 and must next choose a strategy $s_i \in \{P, A, I, V\}$, where

P	= “Paladin”	= {not steal, report},
A	= “Apathetic”	= {not steal, not report},
I	= “Informant”	= {steal, report},
V	= “Villain”	= {steal, not report}.

After choosing their strategies, two of the actors, say i and j , are chosen at random (uniformly) and paired for an interaction. All other actors are *bystanders*. Each bystander receives payoff 1 no matter what others do, yet as bystanders their strategies may affect the payoffs for i and j .

The two selected actors, i and j , are randomly assigned (uniformly) into different roles: one is assigned the first-mover role of “potential criminal,” and the other is assigned the second-mover role of “potential victim.” Their strategies are then carried out in the following way. If the first mover does not steal, then each keeps his or her initial endowment of 1 regardless of the second mover’s report strategy. If the first mover steals, then δ is taken from the second mover, and $\alpha\delta$ (with $0 < \alpha < 1$) is given to the first mover. Hence, an

amount $(1 - \alpha)\delta$ is destroyed (forgone productive activity by the first mover, etc.) in this process, making theft socially inefficient. When the first mover steals, then the second mover’s report-decision plays a role in determining the outcome. If no reporting occurs, final payoffs are $(1 + \alpha\delta, 1 - \delta)$. If reporting occurs, the probability of a conviction is given by $r \equiv (N_P + N_I)/N$ where N_s is the number of actors that chose strategy $s \in \{P, A, I, V\}$. Notice that r , which we shall refer to as the “report rate”, is the proportion of actors that chose report as part of their strategy. If conviction occurs, then the first mover returns δ to the victim and pays an additional punishment cost θ , ending with a payoff of $1 + \alpha\delta - \delta - \theta$. In this case, the victim is fully reimbursed to receive payoff 1. If no conviction occurs, the criminal keeps $1 + \alpha\delta$, and the victim pays an additional cost ϵ (due to a loss in reputation or to retaliation) for a final payoff of $1 - \delta - \epsilon$. Finally, we define $s \equiv (N_I + N_V)/N$, the proportion of actors that chose steal as part of their strategy, as the “steal rate”.

The Paladin, Apathetic, Informant, and Villain labels, which are taken directly from SBD, are meant to convey the essence of the strategies. Paladins act as pure cooperators by not stealing and reporting criminals; Apathetics disengage from society by neither stealing nor reporting (first order cooperator, second order defector); Informants commit crimes but also punish criminals (first order defector, second order cooperator), a strategy that plays a strong role in SBD’s original analysis; and Villains act as pure defectors by both committing and not reporting crimes. We emphasize that these labels do not necessarily match how those terms are used in everyday discourse (if indeed they are used everyday). Rather, they are meant to convey the essence of the four possible strategies in this Adversarial Game. We also note that strategies are chosen at the start of a period and not at information sets within the period. Given the sequential nature of the game, it is natural to consider choices by information set, yet the single choice at the start of the round allows us to consider off-path choices and matches SBD’s original analysis.

SBD’s original Adversarial Game differs from our version in a few ways, two of which are worthy of note and relate to SBD’s primary motivation to examine evolutionary dynamics

of cooperative behavior in adversarial settings. The first difference is that our formulation above is of a one-shot game, whereas their set-up is an evolutionary model with repeated periods. Our one-shot game mimics their stage game in that in each period two actors are selected at random and matched, and all other actors are bystanders for that period. The one-shot game is of sufficient interest and, as we shall see, serves as a useful benchmark for our later evolutionary analysis. The second difference is that strategies in SBD’s version are either inherited or adopted via a pseudo-imitation dynamic (described in detail in Section 4 of this paper) rather than being selected by payoff maximization (best responding). Strategy revision via imitation is commonly assumed in evolutionary settings where the primary interest is the evolution of behavior over time (see Sandholm 2010). We consider a best response dynamic instead of an imitation dynamic because it is a dynamic more closely tied to standard (non-evolutionary) game theoretic analysis and because, as will be shown, it will generate different evolutionary paths of behavior.

The Adversarial Game defined above is a simultaneous game of imperfect information, but it could alternatively be defined as a sequential move game in which the potential criminal makes the steal decision knowing his or her role, and the victim and bystanders make their report decisions after observing a crime. The potential criminal would choose from the strategy set {steal, not steal}, and the others would choose from the set {report, not report}, thus decoupling the steal and report decisions from the overall strategy plan as defined with the simultaneous structure. This sequential structure may more accurately reflect the timing of real-life steal and report decisions; however, we use the simultaneous structure for three reasons. First, having all actors make strategy decisions before knowing their roles makes decisions salient for all actors rather than just those in the matched pair. If bystanders make their report decisions knowing that they are bystanders, then their actions have no relevance to their payoffs. Second, the simultaneous structure better reflects the notion of societal conditions that partly motivates the model. Actors live in a societal setting that has certain properties that pre-date and influence the effectiveness of criminal acts. The

committed behavior of all actors, including bystanders, reflects their latent support, or lack thereof, for punishing criminals that is realized upon the commission of a crime. Third, SBD assume strategy-types are determined at the end of the prior period rather than sequentially within each period. In choosing two-dimensional strategies at the start of the period, we can, when moving to the dynamics, track the full evolution of strategy types, which is a fundamental feature of SBD’s original analysis.

Unlike other social dilemma games, the one-shot Adversarial Game has a single potential deviant. This feature better reflects the inherent asymmetry of criminal behavior, where at a given point in time only one actor may be in a position to take advantage of others or be victimized. The Adversarial Game also formalizes a stylized form of punishment for deviators directly into the basic game by allowing the victim to immediately challenge deviant behavior. Moreover, as with some types of actual criminal behavior, the victim has more to gain than the bystanders when the criminal is punished because he or she is reimbursed. The victim also has even more to lose because an unsuccessful challenge results in an additional private cost. However, the expected result of attempting to punish a deviant depends on the larger societal characteristics, i.e., the behavior of the bystanders. Bystanders, though not directly affected by the realized crime, do indirectly foster or inhibit deviant behavior by influencing the likelihood of successful punishment. Thus, in a very simple way, the Adversary Game captures the positive externalities inherent in the punishment of deviants. The more supportive the environment, the larger the expected benefit of attempting to punish, and the smaller the expected cost. We discuss the importance of this feature of the model in more detail after our analysis of the one-shot game.

3 Static Analysis

3.1 Best-response Functions

An actor’s report decision is only relevant if he or she is a victimized second mover, which occurs with probability $\approx s/N$; with probability $\approx (1 - s)/N$, this report decision is irrele-

vant. Conditional on the decision being relevant, the expected payoff for reporting is higher than when not reporting if

$$\begin{aligned} r(1) + (1-r)(1-\delta-\varepsilon) &> 1-\delta \Rightarrow \\ r &> \frac{\varepsilon}{\varepsilon+\delta} \equiv R. \end{aligned}$$

Hence, the best report decision is to report when $r > R$, not report when $r < R$, and either report or not report when $r = R$.

An actor's steal decision is only relevant if he or she is the first mover, which occurs with probability $1/N$; otherwise, the steal decision is irrelevant. Conditional on the decision being relevant and assuming large N , the expected payoff for stealing is higher than when not stealing if

$$\begin{aligned} r[r(1+\alpha\delta-\delta-\theta) + (1-r)(1+\alpha\delta)] + (1-r)(1+\alpha\delta) &> 1 \Rightarrow \\ r &< \sqrt{\frac{\alpha\delta}{\delta+\theta}} \equiv S. \end{aligned}$$

Hence, the best steal decision is to steal when $r < S$, not steal when $r > S$, and either steal or not steal when $r = S$.

Altogether, the actor's ex ante best response function, assuming $S \neq R$, is thus

$$BR_i(s, r) = \begin{cases} \{P, A\}, & \text{if } s = 0 \text{ and } r > S, \\ \{P, A, I, V\}, & \text{if } s = 0 \text{ and } r = S, \\ \{I, V\}, & \text{if } s = 0 \text{ and } r < S, \\ P, & \text{if } s > 0 \text{ and } r > S \text{ and } r > R, \\ A, & \text{if } s > 0 \text{ and } S < r < R, \\ I, & \text{if } s > 0 \text{ and } R < r < S, \\ V, & \text{if } s > 0 \text{ and } r < S \text{ and } r < R, \\ \{P, A\}, & \text{if } s > 0 \text{ and } S < r = R, \\ \{P, I\}, & \text{if } s > 0 \text{ and } R < r = S, \\ \{A, V\}, & \text{if } s > 0 \text{ and } r = S < R, \\ \{I, V\}, & \text{if } s > 0 \text{ and } r = R < S. \end{cases}$$

This function is depicted graphically in Figure 1.

3.2 Rationalizable Strategies

We use this best response function to obtain our first result.

Proposition 1 *Fix δ , α , ε , and θ . Then, every pure strategy is rationalizable.*

Proof. Consider strategy P . P is a best response to a conjecture that all others choose P , which is a best response for each actor to the conjecture that all others choose P , and so on. An infinite chain of justification can thus be created in which each actor believes all others choose P at each step, thus making P rationalizable.

Next consider strategy A . A is a best response to a conjecture that all others choose P , which we know is rationalizable from above. An infinite chain of justification can thus be created for A .

Now consider strategy V . V is a best response to a conjecture that all others choose V , which is a best response for each actor to the conjecture that all others choose V , and so on. An infinite chain of justification can thus be created in which each actor believes all others choose V at each step, thus making V rationalizable.

Finally consider strategy I . I is a best response if all others choose A , which we know is rationalizable from above. An infinite chain of justification can thus be created for I . ■

As evident by this result, common knowledge of rationality does not alone restrict the set of potentially observable behaviors.

3.3 Perfect Bayesian Equilibrium

The standard solution concept for (Bayesian) games in which the actors act simultaneously not knowing the move by Nature is the Bayesian Nash Equilibrium (BNE) concept. Our second result identifies the BNE of the Adversarial Game.

Proposition 2 *Fix δ , α , ε , and θ .*

(a) *The set of pure Bayesian Nash Equilibria consists of:*

(i) *symmetric profile “Utopia”, in which each actor chooses P ;*

(ii) *symmetric profile “Dystopia”, in which each actor chooses V ;*

(iii) *any asymmetric “Semi-Utopia” profile, in which a fraction z of actors choose P , the remaining fraction $(1 - z)$ of actors choose A , and $S \leq z < 1$;*




Figure 1: A graphical representation of the best response function $BR_i(s, r)$ under the three cases: (a) $s = 0$, (b) $s > 0$, $S > R$, and (c) $s > 0$, $R > S$.

(iv) and the asymmetric “Semi-Dystopia” profile, in which the fraction R of actors choose I and the remaining fraction $(1 - R)$ of actors choose V .

(b) The sum of expected utilities is maximized in Utopia and Semi-Utopia.

Proof. (a) 1. We first show that the profiles listed are equilibria.

(a-i) Utopia with every actor choosing P implies $(s = 0, r = 1)$. From the best response function, we see that P is a best response to $(s = 0, r = 1)$ for each actor. Hence, Utopia is a pure, symmetric BNE.

(a-ii) Dystopia with every actor choosing V implies $(s = 1, r = 0)$. From the best response function, we see that V is a best response to $(s = 1, r = 0)$ for each actor. Hence, Dystopia is a pure, symmetric BNE.

(a-iii) Consider Semi-Utopia with $S \leq z < 1$. It follows that $(s = 0, r = z \geq S)$. From the best response function, we see that both P and A are best responses. With each actor choosing a best response, it follows that Semi-Utopia is an asymmetric, pure BNE.

(a-iv) Consider Semi-Dystopia with the fraction R of players choose I and the fraction $(1 - R)$ of players choose V , so that $(s = 1, r = R)$. From the best response function, each Informant and Villain is playing a pure best response. Hence, this Semi-Dystopia is an asymmetric, pure BNE.

2. It is straightforward to show that for every other (s, r) combination, at least one actor is strictly better off in expectation by changing strategy. The cases left to consider are:

(i) $(s = 0, r < S)$, where non-criminals are strictly better off in expectation by stealing;

(ii) $(s > 0, r > S)$, where criminals are strictly better off in expectation by not stealing;

(iii) $(s > 0, R < r \leq S)$, where non-reporters are strictly better off in expectation by reporting;

(iv) $(s > 0, 0 < r < R)$, where reporters are strictly better off in expectation not reporting;

(v) $(0 < s < 1, r = 0 \text{ or } r = R < S)$, where non-criminals are strictly better off in expectation by stealing.

(b) Social utility is lost anytime a crime occurs, and crime occurs in expectation if and only if at least one actor selects a V or I strategy. Hence, Utopia and Semi-Utopia have maximized sum of utilities. ■

The Utopia and Dystopia labels are taken from SBD, though the social configurations they apply these labels to are slightly different. SBD use Utopia to refer to any state in which no crime occurs, thus combining our definition of Utopia as well as our definition of Semi-Utopia under one label. They refer to both Utopia and Semi-Utopia as Utopia because the distinction was not important given their dynamic analysis. However, our analysis below finds the distinction to be important. SBD use Dystopia to refer to the state in which no punishing occurs, as we do here. However, the SBD Dystopian state includes both Villains

and Apathetics, while ours includes only Villains. There is no analogue of our Semi-Dystopia in the SBD dynamics.

Two technical matters regarding Proposition 2(b) are of note. First, the sum of utilities is maximized in any setting with no crime on the equilibrium path, so it is possible that the sum of *realized* utilities can be maximized with some actors choosing I or V . This occurs when none of those crime committing actors is selected to be the first mover. However, whenever I or V are chosen by at least one actor and before roles are realized, there is a non-zero probability that an inefficient outcome will occur. Second, if a fraction z of actors choose P and the remaining fraction $(1 - z)$ choose A , then the sum of utilities will be maximized with any z , $0 \leq z \leq 1$. In effect, there are efficient, non-equilibrium Semi-Utopia strategy profiles when $0 \leq z < R$.

3.4 Equilibrium Refinements

As seen above, the Adversarial Game has multiple equilibria. Without any further assumptions about how strategies are selected, it is not clear which, if any, of the equilibria would be chosen. One way to approach this equilibrium selection problem is to perform an evolutionary analysis. However, before turning to such analysis, we consider equilibrium refinements in the one-shot game, which can identify which equilibria are most likely to be played and thereby identify which equilibria we expect may arise in our evolutionary analysis.

Although we find many asymmetric Semi-Utopia to be equilibria in the one-shot game, we do not find many asymmetric Semi-Dystopia equilibria. The reason for this result is that the symmetric Utopia with all Paladins is a weak Bayesian Nash Equilibrium while the symmetric Dystopia with all Villains is a strict Bayesian Nash Equilibrium. To see this, observe that if all others are Paladins, then i is indifferent between being a Paladin and an Apathetic, while if all others are Villains, then i 's unique best response is to be a Villain. This fact has implications for refining the set of equilibria.

We consider here two refinements. The first is the notion of an Evolutionarily Stable

Strategy (ESS), which is a strategy that, if adopted by all actors, cannot be invaded by a mutant (Sandholm 2010). Recognizing that the BNE of the Adversarial Game is effectively a Nash Equilibrium, we here consider as an Evolutionarily Stable Strategy in this Bayesian Game a strategy whose expected utilities satisfy the utility conditions for ESS. The second refinement is Trembling Hand Perfect Equilibrium (THPE) in which no actor wants to change his or her strategy given the small chance that others will deviate from their intended strategies (Fudenberg and Levine 1996). As with ESS, we apply the THPE concept using expected payoffs rather than realized payoffs.

Proposition 3 *Fix δ , α , ε , and θ .*

- (a) *V is the only pure Evolutionarily Stable Strategy.*
- (b) *Dystopia and Utopia are the only pure Trembling Hand Perfect Equilibria.*

Proof. (a) Let $U_i(x, y)$ be the expected payoff to actor i when he or she chooses x and all other $N-1$ actors choose y . By definition, a pure ESS is a symmetric pure equilibria in which all actors choose strategy x such that (i) $U_i(x, x) \geq U_i(y, x)$ and (ii) if $U_i(x, x) = U_i(y, x)$ then $U_i(x, y) > U_i(y, y)$, for any strategy $y \neq x$. From Proposition 2(a), the only two symmetric equilibria that may potentially be ESS are all choose V (Dystopia) and all choose P (Utopia).

Consider Utopia. From the best response function it is evident that $U_i(P, P) = U_i(A, P)$ but $U_i(P, A) = U_i(A, A)$. Hence, P is not an ESS. Now consider Dystopia. From the best response function, it is clear that $U_i(V, V) > U_i(y, V)$ for $y \in \{P, A, I\}$. Hence, V (Dystopia) is an ESS.

(b) A THPE is a Nash Equilibrium with certain properties, one of those being that no weakly-dominated pure strategy can be played in a THPE (CITE). From Proposition 2(a), the only equilibria to consider are Dystopia, Utopia, and Semi-Utopia. From the best response function: V is the unique best response when all others choose V , so Dystopia is a THPE; P is the unique best response when s is close to 0 and r is close to 1, so Utopia is a THPE; but A is weakly dominated by P in Semi-Utopia, so Semi-Utopia is not a THPE. ■

That Dystopia withstands typical revisions follows from it being a strict equilibrium. Conversely, all-Paladin Utopia cannot withstand invasions by Apathetics, but it can withstand small mutations or errors in which crime occurs.

3.5 Other Social Dilemma Games and the Reporting Externality

It is readily apparent from Propositions 1 and 2 that one fundamental difference between the Adversarial Game and standard social dilemma games (Prisoners Dilemma, Public Good, Common Property Resource) is that whereas a social dilemma game typically has a single equilibrium that is inefficient, the Adversarial Game has multiple equilibria, one of which is efficient. Indeed, upon closer inspection, though the Adversarial Game is a social dilemma game in spirit and purpose, it is actually an N -actor, four-strategy coordination game with efficient Utopia and inefficient Dystopia as the two focal equilibria. The Adversarial Game is transformed from a social dilemma into a non-standard coordination game via the reporting externality which makes punishment (reporting) a best response when enough others also punish (we refer to it as a non-standard coordination game because coordination is successful on only two of the four strategies). The second order punishment problem is not a typical social dilemma but rather a coordination game that has enough force to transform the first order social dilemma, indeed, the entire game, into a coordination game.

It is instructive to revisit SBD's interpretation of the model with this new insight in mind. SBD explain that the reporting reflects a willingness to cooperate with authorities. That is, the setting is one in which there already exists in place a third-party institutional framework that can leverage society members' willingness to cooperate into effective punishment of criminal behavior. The larger lesson is that if such institutions can be developed, then cooperation can be sustained even in short horizon interactions. This begs the question of how such institutions can be developed in the first place; however, it is evident that such institutions fundamentally change the incentives of potential criminals when there is sufficient societal support.

		2			
		C	D		
1	C	2,2,0	0,3,0	1	C
	D	3,0,0	1,1,0*		D
(A) 3 chooses $r = 0$				(B) 3 chooses $r = 1$	

Figure 2: Modified prisoner's dilemma game.

In principle, any social dilemma game can be converted from a social dilemma to a modified coordination game through a similar mechanism. Consider the following example. Suppose the Adversarial Game payoffs were replaced with a simultaneous move Prisoners Dilemma Game payoffs and then appropriately modified. Also suppose that whether defectors are reported and punished depends only on the choice of a single bystander, actor 3, whose actions represent society's support or lack thereof for punishing defectors. The bystander can choose the reporting rate to be either $r = 0$ or $r = 1$, and his or her payoff does not depend on this choice. The left matrix in Figure 2 depicts the typical Prisoners Dilemma payoffs when the bystander does not report or punish defectors. The payoffs in the right matrix correspond to when there is full reporting and punishment. The payoffs are calculated by subtracting 3 from each defector and adding 2 to each victim; the defector pays back the 2 lost by the defection, which goes to the victim, and then pays an additional 1 as punishment.

In this modified Prisoners Dilemma Game, as with the Adversarial Game, we have defection as the unique best response when there is no reporting, and we have cooperation as the unique best response with full reporting. The two equilibria of Dystopia and Utopia (denoted by * in the matrices) are the only two pure Nash Equilibria in the game. More generally, if actors coordinate their reporting, then the incentives to defect will be overcome by the threat of punishment. These two equilibria reveal it to be a game of coordination.

4 Evolutionary Analysis with Best Response Dynamics

SBD's main purpose is to examine the evolution of behavior with long-run repetition of the Adversarial Game. They assume that strategies switch according to a special imitation dynamic. Say that an actor is a "loser" in a period if that actor's payoff is strictly less than the initial endowment. Clearly, a first mover will only be a loser if he or she was an Informant or Villain that was convicted after being reported against, while a second mover is a loser if he or she is a victimized Apathetic or Villain or a victimized Paladin or Informant that unsuccessfully challenged after being victimized. The "winner" is the other player in that period. SBD assume that only losers switch strategies and that they do so by (possibly imperfectly) mimicking one of the players in that round. The choice of which player to mimic is made with a probability proportional to the players' payoffs for that round. This probability implies a certain amount of inertia in the system. However, a caveat is that any loser who chooses to mimic the second player always becomes a non-criminal type, directly mimicking the second player's reporting strategy only. The idea behind this is that the loser in this case has decided not to mimic the criminal (which would certainly cause him or her to become a criminal type), and has therefore implicitly decided not to commit criminal acts. In short, the SBD dynamic is best described as a modified imitation dynamic with inertia.

SBD show that, with a deterministic version of their imitation dynamic, the Adversarial Game always converges to either Utopia/Semi-Utopia or their form of Dystopia (consisting of mostly Villains with some Apathetics, due to the imperfect mimicking), and that whether the system converges to one or the other depends strongly on the presence of Informants in the initial population. If there are no Informants in the initial population, then the system converges to Dystopia unless a large fraction of the initial population are Paladins; if there are any Informants, then the system converges to Utopia/Semi-Utopia. With a stochastic imitation dynamic, the system converges to Utopia/Semi-Utopia with quickly increasing probability as the initial number of Informants increases. In short, increasing the second

order cooperation (punishment of defectors) among defectors themselves has a powerful effect on the system’s resting state, so much so that any number of initial Informants is sufficient to eventually bring the system to Utopia.

It remains to be seen whether this striking result holds for other strategy revision dynamics. We here consider the best response dynamic, which is an important alternative to the imitation dynamic (Sandholm 2010). In an imitation dynamic, the actor adopts the strategy of another actor, usually with the likelihood of adoption increasing in the performance of the actor potentially copied. Inertia takes two forms: only one or a small number of actors are allowed to switch strategies in a given round, and each of those potential switchers will switch with probability less than 1. In a best response dynamic, the actor that switches does so with more foresight: he or she switches to a strategy that is a best response to the current strategies of the other actors. As with an imitation dynamic, a best response dynamic usually incorporates inertia, with only one or a few actors allowed to switch in a given period.

There are many forms that a best response dynamic can take. Three main characteristics are the rate of inertia, how choice is made among competing best responses, and whether decision errors are allowed during strategy switches. We consider a Best Response Dynamic, denoted SBD-BRD, that selects on these three characteristics to maintain a closeness in spirit to the SBD imitation dynamic. The intent is to find the closest best response dynamic analog to the SBD imitation dynamic, thus allowing for the sharpest comparison of results when switching the dynamic. We then consider an important variation on the dynamic to check robustness.

Definition 1 ***SBD-BRD** is the following switching protocol:*

- (a) *At the end of period t , with probability $q > 0$ one of the two selected actors is chosen to switch strategy.*
- (b) *Conditional on being selected to switch, the actor switches to a best response to the population strategy profile of period t , and if more than one best response exists then one is*

selected at random (uniformly).

Note how SBD-BRD handles inertia, multiple best responses, and switching errors. It intentionally mimics the degree of inertia in SBD's imitation dynamic: there is a chance no actor switches strategy; at most one actor switches strategy; and that actor is one of the two matched actors. SBD's assumption that the loser is the one to switch matches the spirit of imitation dynamics because a loser would want to imitate a winner but not vice versa. However, in the spirit of the best response dynamic, potentially any actor may see benefits in switching, so we allow either of the two actors to be the one to switch. SBD's imitation dynamic does not have a natural analog to the issue of multiple best responses in the best response dynamic, so without guidance from the SBD dynamic we here adopt a conventional assumption that the actor mixes equally among best responses. SBD's imitation dynamic also does not include switching errors, and SBD-BRD makes a similar assumption.

Proposition 4 *Fix δ , α , ε , and θ . Under SBD-BRD, there are two possible long-term behaviors:*

- (a) *the system evolves to Dystopia, or*
- (b) *the system maintains a level of reporting $r \gtrsim S > R$.*

Proof. Here, we assume that $|R - S| \gg 1/N$, and define the following regimes of (s, r) space, which cover all possible values:

- (i) $s > 0$, $r < R$,
- (ii) $s = 0$, $r \geq S$,
- (iii) $s = 0$, $r < S$,
- (iv) $s > 0$, $r \geq R$.

(a) First, consider the case $S < R$, with initial conditions in regime (i). Here, the best response is to not report, so that each strategy update will either leave r unchanged or will decrease r . There is a nonzero probability that these sequences of updates will also leave $s > 0$. Conditional on s remaining nonzero, then, the system is guaranteed to eventually

evolve to a point in which $r < S$, in which case V is the unique best response, leading the system to eventually reach Dystopia and stay there forever. If s becomes zero before this can happen, the system switches to regime (ii).

Now, consider the case $S < R$ with initial conditions in regime (ii). If $r > S$, either P or A is a best response, so each update will keep $s = 0$ and may increase, decrease, or maintain r . With probability 1, then, the system will eventually evolve through neutral drift to either a point in regime (iii) or to the specific point $s = 0$, $r = S$. If $r = S$, all strategies are best responses, so eventually an update will occur that either creates a criminal and puts the system back in regime (i), maintains $s = 0$ and causes $r > S$ (still in regime (ii)), or puts the system into regime (iii). Hence, the eventual end result of beginning in regime (ii) is to either enter regime (i) or regime (iii).

Let us consider, then, $S < R$ with initial conditions in regime (iii). Here, the best responses are V and I , so the next update will necessarily put the system into regime (i) above.

Next, consider the case $S < R$, with initial conditions in regime (iv). If $r = R$, either P or A is a best response, so there is a nonzero probability that r decreases below R , leading to regime (i). Otherwise, the unique best response is P , causing r to either increase or stay the same each round, and causing s to decrease or stay the same each round. Eventually, then, the system will evolve to a state in which $s = 0$, $r > R > S$, which is regime (ii).

Hence, when $S < R$, initial conditions in regimes (ii)-(iv) will all eventually (though possibly indirectly) lead to regime (i), which leads to the absorbing Dystopia state with a nonzero probability, and otherwise leads back to regime (ii). Therefore, when $S < R$ the system will eventually evolve into, and forever remain in, the Dystopian state, regardless of initial conditions.

Finally, consider the case $R < S$ with initial conditions in regime (i). The best response here is V , causing the system to evolve to Dystopia with probability 1.

(b) Consider now $R < S$, with initial conditions in regime (iv). If $r = R$, both I and

V are best responses, so the system may evolve to regime (i) and, therefore, to Dystopia. However, there is a nonzero chance that the system instead evolves such that $r > R$. With r strictly greater than R , reporting is always the best response, so r will only increase or maintain. Furthermore, if $r < S$, we are guaranteed to maintain $s > 0$, as I is the unique best response. Therefore, the system will eventually evolve to a point with $s > 0$, $r \geq S$. At this point, P is certainly a possible best response (usually the unique best response), so the system will eventually evolve into regime (ii).

Initial conditions in regime (ii) behave the same way when $R < S$ as when $R > S$, though this behavior may lead to entering different subsequent regimes, due to the switching of the order of R and S . That is, initial conditions (ii) will eventually lead through neutral drift either to regime (iv), specifically with $S - \frac{1}{N} \leq r$, or regime (iii), specifically with $S - \frac{1}{N} \leq r < S$.

Now we turn to $R < S$ with initial conditions in regime (iii). Here, the best response is I or V , so the next update will cause us to leave regime (iii) as s becomes nonzero. But, there are three possibilities for how this will happen. First, if $r \geq R + \frac{1}{N}$, this update will put us into regime (iv). This is guaranteed to happen if we entered regime (iii) from regime (ii), as mentioned above, and will specifically lead to regime (iv) with $r \geq S - \frac{2}{N}$. If $r < R - \frac{1}{N}$, this update will put us into regime (i), and lead to Dystopia. If $R - \frac{1}{N} \leq r < R + \frac{1}{N}$, the system may evolve into either (iv) or (i).

We can therefore summarize these results in the following way. For $R < S$, if the system is ever in a state in which $r \geq S - \frac{2}{N}$, r is guaranteed never to fall out of that region again. In this case, the system will maintain an r in this range while cycling between regimes (iv) and (ii) (potentially with very brief stops in (iii) between). Furthermore, once this cycling begins, there will never be more than 1 Villain in the system at a time, and never more than $SN + 1$ Informants (the typical number will be much less than this). Finally, the only initial conditions that are guaranteed to not lead to this long term behavior (and are guaranteed to enter Dystopia) are those in regime (i) or those in the subset of regime (iii)

with $r < R - \frac{1}{N}$; essentially those initial conditions with $r \lesssim R$. All other initial conditions are either guaranteed to end in this cycling (when $r \gtrsim R$) or may lead to either this cycling or to Dystopia (when $r \approx R$). ■

The above argument is explored graphically in Figure 3. Here, we have plotted the mean flow fields in (r, s) space for the two cases $R > S$ and $S > R$, using $N = 20$. These mean fields assume a uniform distribution on the initial number of players of each strategy. The background colors range from light gray to bright orange, and represent the total amount of time spent in each of the available cells, from very little to very much. For the case $R > S$ (left), we see that the trajectories are separated along the line $r = R$ when $s > 0$, but that the $s = 0$ line serves as a funnel from the region $r > R$ to the region $r < R$. When $r < R$, the trajectories inevitably lead to Dystopia, where they remain forever (hence the brightest orange square on the figure). For the case $S > R$, we also see a separation around $r = R$, but the line $s = 0$ no longer funnels trajectories from above this separation to below it. Here, we have treated $r = S$, $s = 0$ as an absorbing point (hence the brightest orange on that plot), because any trajectory passing through it will in fact never settle down to a single point, and will forever cycle near that region.

Proposition 4 reveals that the SBD result about the influential role of Informants does generalize in the case of SBD-BRD, as long as $R < S$, though the role that the Informants play differs from SBD to best response. In SBD, Informants are a sufficient but not necessary condition to drive the system to Utopia/Semi-Utopia. In SBD-BRD, on the other hand, Informants are a necessary but not sufficient condition to keep the system near Utopia/Semi-Utopia. This result is formalized in the following proposition, which considers a modified form of SBD-BRD.

Proposition 5 *Fix δ , α , ε , and θ . Under SBD-BRD, but with the strategy I now disallowed, the system will always evolve to Dystopia in the long term.*

Proof. With I disallowed, we must now reconsider our Best Response function $BR_i(s, r)$. For all entries that list I as a non-unique best response, we may simply remove I . This




Figure 3: Mean flow fields in (r, s) space for the two cases $R > S$ (left) and $S > R$ (right), using $N = 20$. Colors represent the total amount of time spent at a cell, ranging from light gray (little time) to bright orange (much time), assuming a uniform initial distribution on the number of players of each strategy type. For $R > S$, all trajectories eventually end in Dystopia. For $S > R$, we have treated $r = S$ as an absorbing point (all trajectories passing through that point actually cycle near it indefinitely), so that all trajectories end either there or at Dystopia.

leaves only the special case $R < r < S$, for which I is the unique best response. Note, therefore, that if $S < R$, removing I from our possible list of strategies does not change the Best Response function in any qualitative way, so that the result above – that if $S < R$ the system will always evolve to Dystopia – still holds. Also, note that if I is disallowed, the inequality $s \leq 1 - r$ must hold, since in this case $s = V/N$ and $r = P/N$, and $V = N - A - P$.

If $R < S$, we must determine the new best response for $R < r < S$. Since $r < S$, V is superior to A , and since $R < r$, P is superior to A . Hence, either V or P is the best response, or both. The expected payoffs for P and V are

$$1 - \frac{s}{2}(1 - r)(\delta + \varepsilon) \text{ and } 1 - \frac{s}{2}\delta + \frac{1}{2} [\alpha\delta - r^2(\delta + \theta)] ,$$

respectively. These payoffs are equal when $s = Q(r)$, with

$$Q(r) = \frac{(\delta + \theta)(S^2 - r^2)}{(\delta + \varepsilon)(r - R)} .$$

Thus, if $s > Q(r)$, P is the best response, if $s < Q(r)$, V is the best response, and if $s = Q(r)$, both P and V are the best response.

The function $Q(r)$ is defined on $s \in (R, S]$, is decreasing and concave up with slope guaranteed to satisfy $Q'(r) \leq -2$, and $Q(S) = 0$ while $Q(s)$ diverges as $r \rightarrow R$. Furthermore, there is a critical value of r , call it r_* , where the curve $s = Q(r)$ intersects the line $s = 1 - r$ (the maximum value s can attain for any given r), and $R < r_* < S$. Therefore, V is always the unique best response if $r < r_*$ since in this case s cannot possibly be greater than $Q(r)$. We now introduce two new regimes for initial conditions, based upon $Q(s)$:

(v) $s > 0$, $r \geq Q(s)$,

(vi) $s > 0$, $r < Q(s)$.

Now, consider $R < S$ with initial conditions in regime (v). If $r = Q(s)$, the next update may either remain in regime (v) with $r > Q(s)$, put the system into regime (vi), or put the system into regime (ii). For $r > Q(s)$, P is the unique best response, so the system will eventually evolve to a state in regime (ii).

Consider $R < S$ with initial conditions in regime (ii). Here, through neutral drift the system will eventually evolve into one of three other regimes: regime (v), specifically with $s = \frac{1}{N}$, $r = S$; regime (iii), specifically with $S - \frac{1}{N} \leq r < S$; or regime (vi), specifically with $s = \frac{1}{N}$, $r = S - \frac{1}{N}$. To see this last possibility, it is important to remember that $Q'(r) \leq -2$ and $Q(S) = 0$, so that the point specified certainly resides within regime (vi) and not regime (v).

Consider next $R < S$ with initial conditions in regime (iii). Here, the unique best response is V , so the next step will bring the system out of this regime and into either regime (vi), which is always a possibility, or regime (v), which is only a possibility if $S - \frac{1}{2N} \leq r < S$. As above, this limitation on the possibility of bringing the system into regime (v) relies on the facts that $Q'(r) \leq -2$ and $Q(S) = 0$.

Finally, consider $R < S$ with initial conditions in regime (vi). Here, V is the unique best response, so each update that changes r or s is guaranteed to either decrease r while increasing s (if a Paladin becomes a Villain) or maintain r while increasing s (if an Apathetic becomes a Villain). If $r \geq r_*$, then there is a nonzero probability that the system will undergo

a sufficient number of these latter updates to bring it back into regime (v). But, there is also a nonzero probability that a sequence of updates occurs whereby the system remains in regime (vi) and attains $r < r_*$. Below this point, regime (v) can never be re-entered, so the system evolves to Dystopia.

Hence, if $R < S$ and I is disallowed, all initial conditions will eventually bring the system into regime (vi), which leads to Dystopia with a nonzero probability. Therefore, all initial conditions are guaranteed to eventually evolve to Dystopia, where the system will then remain forever. ■

5 Conclusion

SBD’s evolutionary game showed that Informants alone may drive a society from Dystopia to Semi-Utopia. As a consequence, its primary policy implication is that recruiting Informants from the general population may be an overall valid approach in helping reduce crime. SBD’s result stems from the primary role Informants play in their imitation dynamics, where Informants render Dystopia unstable to small perturbations. In SBD’s case, then, Dystopia may be reached under a set of initial conditions containing no Informants, but any infinitesimal deviation from this will drive the system to Semi-Utopia in the long run. This is a universal result, in the sense that it holds across parameter space, as long as Informants are present. It becomes natural then to ask what is the best way to recruit Informants from the general population – and hasten the transition towards Semi-Utopia – given specific costs associated with different player’s current strategies and histories. This question is explored in Short, Pitcher, and D’Orsogna (2012).

Our current analysis presents a more nuanced picture: we find that under the best response dynamics, the role Informants play depends on parameter choices, and that their mere presence does not necessarily drive the system to Utopia as they did in SBD’s original work. In particular, when $S < R$, Informants do not play a central role because the strategy of committing crimes and reporting to authorities is never a best response. On the other

hand, when $R < S$, the Informant strategy plays a pivotal role in evolving the system toward, and then maintaining, a level of reporting $r \gtrsim S$, as long as the initial reporting level $r_0 \gtrsim R$. Thus, the availability of I as a possible strategy is the only channel allowing the system to reach and maintain a state very near Semi-Utopia, making Informants a necessary, but not sufficient, criteria for the emergence of a low-crime state.

As discussed earlier, simply converting citizens to Informants in the context of best response dynamics will not necessarily guarantee the system to evolve towards Utopia as it did in SBD's dynamics. Carefully choosing parameters, however, may allow the system to make this transition; specifically, if parameters are chosen so that $R < S$. In terms of the original model parameters, the $R < S$ constraint implies that several conditions must be met in order for the low-crime state to emerge. Let us first note that we may assume α , δ , and ε to be fixed parameters intrinsic to the game between victim and victimizer, that cannot be adjusted by law enforcement. We may instead allow the degree of punishment θ to be a variable parameter that authorities may fine tune at will. In terms of θ , then, the constraint $R < S$ implies that

$$\theta < \frac{\delta}{\varepsilon^2} [\alpha(\varepsilon + \delta)^2 - \varepsilon^2] . \quad (1)$$

In other words, the degree of punishment imparted on criminals must not be too onerous, so as not to discourage the strategy I from being chosen. Hence, one policy implication of the current work is that punishments should not be made too harsh, lest Semi-Utopia be made unobtainable.

Furthermore, note that Eq. 1 can be made true only if $\alpha > \varepsilon^2/(\varepsilon + \delta)^2$. This implies, though, that for a given $\alpha < 1$, criminals may always choose a set of parameters δ and ε that violates this second inequality, by choosing

$$\varepsilon \geq \frac{\delta\sqrt{\alpha}}{1 - \sqrt{\alpha}}, \quad \delta \leq 1 - \sqrt{\alpha} ; \quad (2)$$

the latter condition arises from the constraint that $\delta + \varepsilon \leq 1$. Thus, if retaliation against witnesses is too large, the system has no chance of reaching a low-crime state, regardless

of θ . Of course, this raises another game theoretic challenge, this time for criminals: to guarantee a crime-friendly society, criminals must restrain themselves to stealing no more than $\delta = 1 - \sqrt{\alpha}$, but each criminal is individually motivated to choose a δ as large as possible, i.e., 1. Hence, it seems plausible that only in societies where criminals themselves are highly organized can such a situation arise, and if each criminal is acting purely in his own best interest, the authorities should be able to arrange for punishment levels that promote a Semi-Utopian state.

Of course, whether an imitation, best response, or some other dynamic best captures actual decisional strategies is an empirical question. Are players perfectly rational and capable of choosing a best response in each one-shot game, perhaps allowing for an initial learning phase, or do they rather learn by imitation and by adaptation? It would thus be interesting to verify in actual subjects whether players follow either of these two dynamics – best response or imitation – or some other decisional process when faced with the choice of committing and reporting crimes. We leave this as future work.

References

- [1] J. Beittel. 2009. *Mexico's Drug Related Violence*. Washington, D.C.: Congressional Research Service.
- [2] S. Bowles, H. Gintis. 2011. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton, NJ: Princeton University Press.
- [3] R. Boyd, H. Gintis, S. Bowles. 2010. "Coordinated Punishment of Defectors Sustains Cooperation and Can Proliferate when Rare." *Science* 328: 617-620.
- [4] R. Bursik, H. Grasmick. 1993. *Neighborhoods and Crime: Dimensions of Effective Community Control*. New York, NY: Lexington Books.
- [5] J. Carpenter. 2007. "Punishing Free-riders: How Group Size Affects Mutual Monitoring and the Provision of Public Goods." *Games and Economic Behavior* 60: 31-51.
- [6] A. Chaudhuri. 2011. "Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature." *Experimental Economics* 14: 47-83.
- [7] M. Cinyabuguma, T. Page, L. Putterman. 2005. "Cooperation under the Threat of Expulsion in a Public Goods Experiment." *Journal of Public Economics* 89: 1421-1435.
- [8] B.C. Eaton, J. Wen. 2008. "Myopic Deterrence Policies and the Instability of Equilibria." *Journal of Economic Behavior and Organization* 65: 609-624.

- [9] I. Ehrlich. 1996. "Crime, Punishment, and the Market for Offenses." *Journal of Economic Perspectives* 10: 43-67.
- [10] E. Fehr, U. Fischbacher. 2004. "Third-party Punishment and Social Norms." *Evolution & Human Behavior* 25: 63-87.
- [11] E. Fehr, S. Gächter. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review* 90: 980-994.
- [12] J. Fender. 1999. "A General Equilibrium Model of Crime and Punishment." *Journal of Economic Behavior and Organization* 39: 437-453.
- [13] D. Fudenberg, D. Levine. 1996. *Game Theory*. Cambridge, MA: MIT Press.
- [14] D. Fudenberg, P. Pathak. 2010. "Unobserved Punishment Supports Cooperation." *Journal of Public Economics* 94: 78-86.
- [15] D. Gambetta. 1988. *Trust: Making and Breaking Cooperative Relations*. Oxford, UK: Blackwell.
- [16] C. Hauert et al. 2007. "Via Freedom to Coercion: The Emergence of Costly Punishment." *Science* 316: 1905-1907.
- [17] J. Heinrich, R. Boyd. 2001. "Why People Punish Defectors: Weak Conformist Transmission can Stabilize Costly Enforcement of Norms in Cooperative Dilemmas." *Journal of Theoretical Biology* 208: 79-89.
- [18] J. Ledyard. 1995. "Public Goods: Some Experimental Results." In J. Kagel and A. Roth, eds., *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- [19] D. Levine, W. Pesendorfer. 2007. "The Evolution of Cooperation Through Imitation." *Games and Economic Behavior* 58: 293-315.
- [20] E. Ostrom. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York, NY: Cambridge University Press.
- [21] E. Ostrom, J. Walker, R. Gardner. 1992. "Covenants With and Without a Sword: Self-Governance is Possible." *American Political Science Review* 86: 404-417.
- [22] R. Sampson, W. Groves. 1989. "Community Structure and Crime: Testing Social-Disorganization Theory." *American Journal of Sociology* 94: 774-802.
- [23] W. Sandholm. 2010. *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press.
- [24] W. Skogan. 1990. *Disorder and Decline: Crime and the Spiral of Decay in American Neighborhoods*. Los Angeles, CA: University of California Press.
- [25] M. B. Short, P. J. Brantingham, M. R. D'Orsogna. 2010. "Cooperation and Punishment in an Adversarial Game: How Defectors Pave the Way to a Peaceful Society." *Physical Review E* 82, 066114.
- [26] M. B. Short, A. Pitcher, M. R. D'Orsogna. 2012. "External Conversions of Player Strategy in an Evolutionary Game: A Cost-benefit Analysis through Optimal Control." Manuscript.

- [27] S. Takahashi. 2010. "Community Enforcement when Players Observe Partners' Past Play." *Journal of Economic Theory* 145: 42-62.
- [28] E. Xiao, D. Houser. 2011. "Punish in Public." *Journal of Public Economics* 95: 1006-1017.